Assessing Design Tradeoffs in Deploying Undersea Distributed Sensor Networks

Russell Costa and Thomas A. Wettergren

Naval Undersea Warfare Center, Division Newport
Newport, Rhode Island 02841
Email: costar@npt.nuwc.navy.mil
Email: t.a.wettergren@ieee.org

***Abstract*—In this paper we explore design issues in the deployment of distributed sensor networks (DSNs). In particular, the search performance of a notional surveillance network, and its dependence on sensor placement (for a fixed number of sensors) is studied. We describe a search objective for systems of sensors that utilize spatio-temporal techniques to combine individual sensor detections (which are spatially consistent with expected target behavior) in order to determine target existence in the field (track coverage). We utilize this objective in a genetic algorithm to optimize the search coverage as a function of the sensor density within a fixed search region. Target dynamics are treated as parameters with associated probability distributions, and enter the search objective as random parameters. We present several examples of optimal placement given target dynamics and sensor characteristics.**

I. Introduction

Distributed sensor networks have been proposed as viable solutions to a number of problems, including ad hoc communication, environmental monitoring, and various surveillance applications. In the case of underwater surveillance, distributed networks of sensors are of interest to detect and/or track targets for various tactical purposes, one example being coastal security. Such networks are becoming increasingly practical with evolution in sensor technology and reduced costs; however, technical guidance on how best to deploy these systems given the number of sensors to be deployed, the individual sensor characteristics, system fusion criteria, and potential target dynamics, remains a nontrivial technical question.

A significant number of papers have been published related to "coverage" of sensor networks, particularly for wireless communication [1], [2], [3]. The emphasis of "coverage" for these systems, however, is that of area coverage. That is, the objective is to place sensors so that they spatially cover an area, with the emphasis on obtaining information and then passing that information on to some central processing node. Thus optimization problems related to these applications seek to maximize objectives that are information theoretic in nature [4]. The applications for which area coverage is key are typically airborne communications. In the case of undersea surveillance, the expense of underwater communication (sensor to sensor) and the sensors themselves prohibit having the large numbers or large coverage per sensor, which enable the problem to be treated as area coverage. Thus underwater DSN systems should emphasize "track coverage" as a measure of search effectiveness. Track coverage in this context is the ability to detect particular targets, with specific characteristics, both in behavior (motion) and acoustic emission.

In this paper, we consider the optimal placement of a fixed number of sensors, where each sensor employs independent but identical detection criteria, referred to as local detections, and the requirements for combining multiple local detections (fusion), into a decision on the presence (or absence) of a target, are specified. The requirements for deciding if a target is present are met if a predefined number of individual sensor detections are received within a fixed interval of time, where the sensor locations meet spatial requirements kinematically consistent with target motion (the target track). The final determination of a target's presence is not made without meeting, or exceeding, the required number of detections defined by the fusion criteria. Optimality is defined with respect to a search objective, which includes target motion parameters stochastically through probability density functions on speed, course, and location, as well as system and sensor parameters. We use an expression for track coverage which was previously developed for modeling this process [5]. We use this measure of field level performance, referred to as probability of successful search, P_{SS}, as an objective to be maximized by judicious placement of sensors.

II. Search Objective

In order to quantify performance of a distributed sensor network, a modeling framework consisting of sensor characterization (as a function of target and environment), target dynamics, and fusion parameters must be defined. We build the search objective by combining these integral pieces. To formalize the problem to be solved we first define a search region $\mathcal{S} \subset R^2$, which we are interested in for surveillance. The potential target tracks will be defined on (but not restricted to) $\mathcal{S}$. We define $f(\mathbf{z}) : R^2 \to R$ as the underlying sensor density function, and seek the functional form of the sensor density which performs best with respect to our search objective (maximizes the objective).

0-933957-35-1

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **29 SEP 2007**	2. REPORT TYPE	3. DATES COVERED **00-00-2007 to 00-00-2007**

4. TITLE AND SUBTITLE **Assessing Design Tradeoffs in Deploying Undersea Distributed Sensor Networks**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Naval Undersea Warfare Center,Division Newport,Newport,RI,02841**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM002047. Presented at the MTS/IEEE Oceans 2007 Conference held in Vancouver, Canada on Sep 29-Oct 4, 2007. U.S. Government or Federal Purpose Rights License.

14. ABSTRACT
See Report.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**5**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

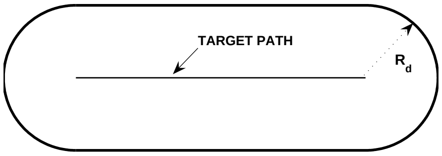

Fig. 1. Geometry of Target Pill Ω_T

A. Modeling Sensor Performance

We adopt a simple parametric form for describing individual sensor performance within the network. Each sensor has an independent detection capability and is represented by a radius of detection, R_D, with a corresponding probability of detecting a (particular) target within this detection radius, P_D. While the detection criteria must be defined to be target specific, both parameters can be represented with varying complexity. For instance, in some cases it may be argued that they be represented by constants (simple), while in other cases it may be necessary to represent the detection parameters with multidimensional functions (complex). The detection radius is defined, such that the integration time of the sensor is accounted for, thus if any part of a target track intersects a sensor circle, that track is detected with associated probability P_D. In this paper we will only be interested in simple parametric forms of sensor detection performance, that is constant P_D for given R_D where R_D could be a function relative to the sensor's position in $\mathcal{S}$.

In distributed sensor networks employing a multiple detection fusion criteria, interaction of an individual sensor and a target track leads to an intuitive geometric description. Figure 1 shows a "target pill" which is the result of a target of fixed velocity, v, over a predefined listen interval, δT, combined with the detection radius of an individual sensor. This is an interpretation common in search theory. The system declares a target present when there are k or more independent sensor detections within this pill-shaped region, where k is a predefined fusion parameter. This region, $\Omega_T \subset \mathcal{S}$, thus represents all possible sensor locations in $\mathcal{S}$, that would detect a particular track. The probability of an individual sensor being in Ω_T is

$$\phi = \int_{\Omega_T(\mathbf{z}_T,\theta,v\delta t)} f(\mathbf{z})\, d\mathbf{z}. \tag{1}$$

where ϕ is a function of the track parameters, velocity (v), location ($\mathbf{z}_T$), and course (θ), as well as the underlying sensor density function, $f(\mathbf{z})$. Note that track location vector $\mathbf{z_T} \subset \mathcal{S}$, while the sensor position vector $\mathbf{z} \subset \mathbf{R^2}$. The probability of successful search (for an individual track) is then

$$P_{ST}(N_D \geq k) = 1 - \exp(-NP_D\phi) \sum_{m=0}^{k-1} \frac{(NP_D\phi)^m}{m!} \tag{2}$$

where N_D is the number of sensor detections within the "target pill", which must meet or exceed the requirement for number of local sensor detections (meeting the fusion requirement) k, for a particular track to be detected, and N is the total number of sensors in the field (fixed). This expression is derived by noting that the track coverage (fusion of multiple detections at the information level) process can be modeled statistically by a binomial process, as the probability of a track being detected requires a combination of k out of N sensors (independent) being in Ω_T with fixed probability ϕ. This binomial process can then be approximated by a Poisson process under the usual assumptions [5].

B. Track Coverage Optimization

The single track detection process is generalized to a field search objective by representing the target dynamic parameters, namely course, speed, and location, by probability density functions (PDFs), $f_\theta(\theta)$, $f_v(v)$, and $f_T(\mathbf{z}_T)$, respectively. From equation (2), the single track successful search probability, and the PDFs of the track parameters, we write the probability of successful search of the field of sensors as

$$P_{SS} = \int_0^{2\pi} \int_{v_{min}}^{v_{max}} \int_{\mathcal{S}} P_{ST}(\phi)\, f_T(\mathbf{z}_T) f_v(v) f_\theta(\theta)\ d\mathbf{z}_T\, dv\, d\theta \tag{3}$$

The dependence of P_{ST} on the track parameters (and the sensor density function) is implicit through its dependence on ϕ. It is seen from this expression that the field measure of search effectiveness is found by marginalization of the track parameters. That is, the uncertainty in the track parameters is "integrated out" in order to calculate the probability of detecting the target using the fusion of multiple detections, as described above.

Equation (3) is the search objective which we optimize to obtain guidance on sensor placement for various tactical scenarios. All parameters are defined as described above with the only unknown being the underlying sensor density function, $f(\mathbf{z})$. Therefore, the optimization problem is to find the sensor density function which results in the best search coverage. To utilize a standard technique for optimization, namely a genetic algorithm, we parameterize the sensor density function as a gaussian mixture. The standard form of a circular gaussian mixture in R^2 is written as

$$f(\mathbf{z}) = \frac{1}{2\pi\sigma^2} \sum_{j=1}^{M} w_j \exp\left(-\frac{1}{2\sigma^2}(\mathbf{z}-\mathbf{z}_j)^T(\mathbf{z}-\mathbf{z}_j)\right). \tag{4}$$

Gaussian mixtures are well-suited to representing unknown smooth functions and can represent to any desired accuracy if enough terms are chosen [6] (i.e. as $M \to \infty$). In practice, we limit our approximation to a reasonable number of mixture terms of $O(100)$. In the parametrization of equation (4), $f(\mathbf{z})$ is represented in terms of the Gaussian variance σ^2, the centers of the mixture components (mean vectors) $\{\mathbf{z}_1, \ldots, \mathbf{z}_M\}$ and the mixture weights $\{w_1, \ldots, w_M\}$. To facilitate the optimization we heuristically select the number of mixture components, M, and distribute the locations on a uniform grid over the search region (i.e. the modes are centered at fixed positions). The variance parameter of each mode is also fixed (and identical for all modes) and is determined to achieve a prescribed level of "flatness" for equal weights. Thus for equal weights we have a reasonable representation of a uniform sensor distribution over the search region $\mathcal{S}$. Fixing these parameters leaves only the weights of the mixture as variable, and thus the weights are the parameters we use to optimize the search objective. Combining equations (3) and (4) into the objective leads to the following optimization problem:

Problem OPT

$$\max_{f(\mathbf{z})} P_{SS} \tag{5}$$

$$\text{subject to} \qquad \sum_{j=1}^{M} w_j = 1, \qquad w_j \geq 0 \text{ for all } j \tag{6}$$

This optimization is performed using a genetic algorithm [7] to search the parameter space (the weights), with each parameter represented by an 8 bit binary string, so that the overall string length in the optimization is $M * 8$. For each generation the mating is performed using single crossover and we always keep as a survivor the string which resulted in the highest value of our objective (P_{SS}). The grid setup maintains good order in our string representation; however, one weakness to this representation is the mapping of the equality constraint of the weights. Weight parameter values are allowed to be in the interval $[0, 1]$ while they must sum to one for the mixture representation to integrate to one. Thus the algorithm normalizes the weights (to sum to one) before calculating the objective for a particular string. This results in string representations that may not be one-to-one mappings into the real valued parameter weights. This can affect convergence; however, even in the presence of this issue, this parametrization leads to reasonably good convergence in our problem as will be shown in the next section.

III. Numerical Examples

We now show some numerical examples which illustrate the utility of optimizing the search objective (3), through parametrization of the sensor density function (4), as a precursor to deploying a field of sensors. In each example the search region $\mathcal{S}$ is a 1 square mile box. We use the search objective to first find an optimal sensor distribution, and then use the resulting distribution to place $N = 50$, passive acoustic sensors. In these examples the listen interval, δT, is 500 seconds, and the required detections, k, are set at 3. The sensor detection probability for the target of interest is a constant $P_D = 0.99$ (cookie cutter) within the detection radius R_D. The target dynamics are represented as a fixed speed of 5 ft/s (approximately 3 knots), and uniform (random) course and position. "Position" refers to the reference position $\mathbf{z}_T$ for each track, which for these examples is the center of the line segment representing the track, defined in $\mathcal{S}$. As the track parameters are only defined in $\mathcal{S}$, the PDFs which represent them must integrate to one over $\mathcal{S}$.

The parametrization of the sensor PDF is represented by $M = 81$ modes (9x9 grid), with $\sigma = 2R_D$. In these examples, we ran 1000 generations with a population size of 40. The selection criteria utilized at each generation was "roulette" [7] which is equivalent to drawing samples from an empirical distribution, based on the normalized objective values of each string in a population. The output of the genetic algorithm is the sensor density function that has performed best (over all the generations) with respect to the search objective. This is not guaranteed to be a local or global maximum, but for our purposes this will represent the optimal distribution with respect to the search objective.

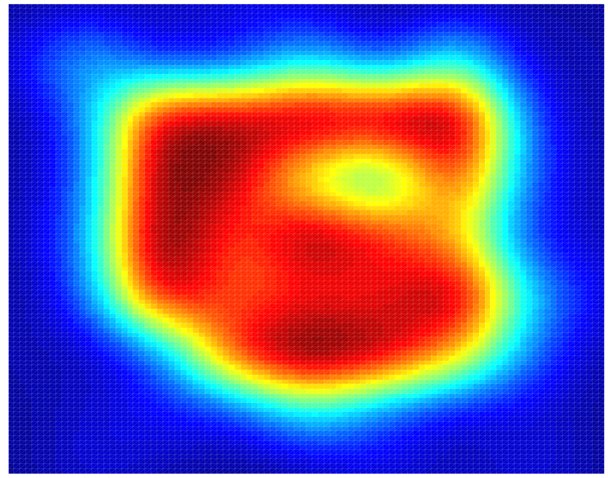

Fig. 2. Genetic Result for Sensor PDF for Uniform Course (k=3)

A. Example 1: Uniform Detection Radius

In this example we take the sensor detection radius to be fixed over $\mathcal{S}$ at $R_D = 200$ feet. The algorithm was started with uniform weights which corresponded to a P_{SS} of 0.1762 and achieved a P_{SS} of 0.3067 as a result of the genetic search. Figure 2 shows the sensor density (restricted to $\mathcal{S}$) after running the genetic algorithm for 1000 generations. The interpretation of the color map is that red represents highest probability of a sensor being, while blue is the lowest. In this example, the sensor density has increased mass toward the center, while the mass along the search region boundary is reduced. This is a result of the multiple detection requirement, $k = 3$, and the small number of sensors deployed $N = 50$. In order to meet the fusion criteria the sensors need to be packed in toward the middle to maximize track coverage within the region. Note that although the optimization is of the sensor density function, the search objective, P_{SS} is a function of N, the fixed number of sensors to be deployed in $\mathcal{S}$.

The objective values P_{SS}, shown above, are with respect to a sensor distribution defined on R^2, the probability of more practical interest is one normalized to $\mathcal{S}$ (since for practical problems we will only place sensors in the search region). This can be calculated by sampling from the empirical distribution, formed from the genetic result by evaluating the sensor density on a fixed grid (discrete) and normalizing the mass function to sum to 1 on $\mathcal{S}$. These cells are then sorted by weight,

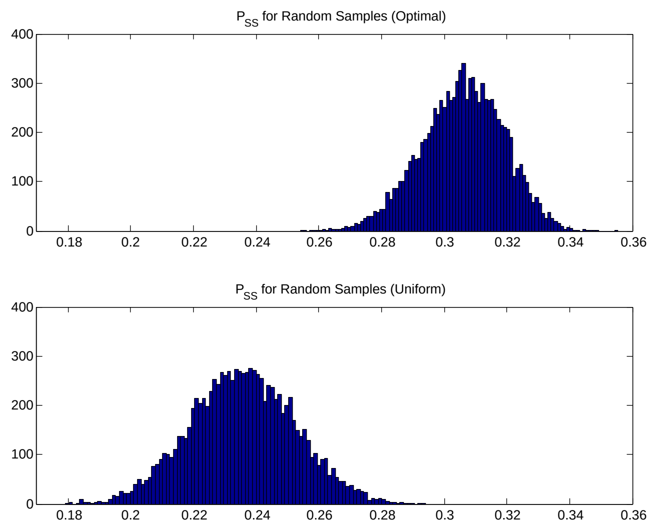

Fig. 3. Example 1 P_{SS} Sampling Results

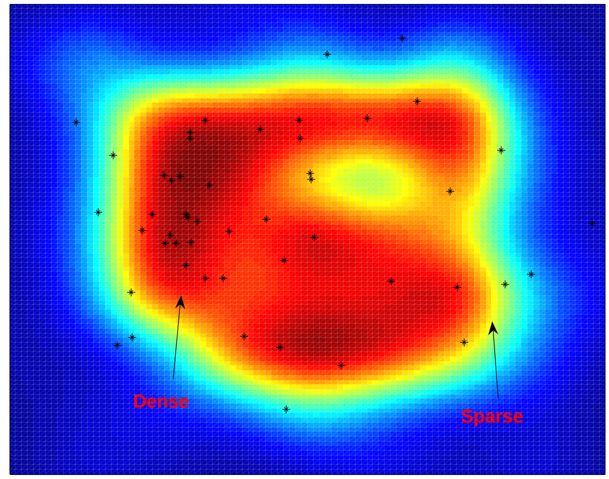

Fig. 5. Example 1 Random Sample (min P_{SS})

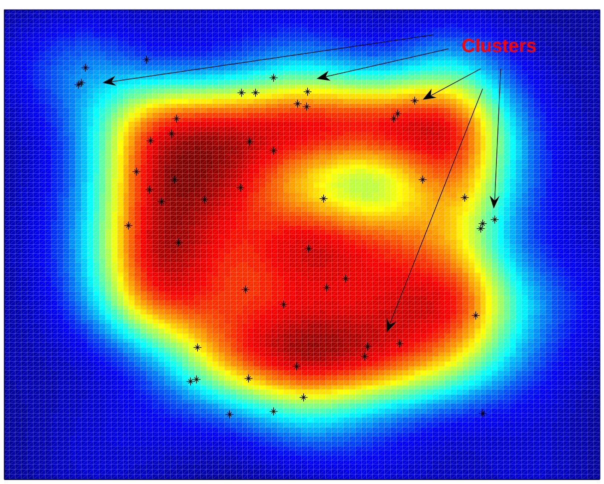

Fig. 4. Example 1 Random Sample (max P_{SS})

and cells are selected (sampled) by mapping random uniform numbers into intervals that are proportional to the weight of each cell. As the number of cells goes to infinity this empirical representation converges to the actual distribution. Thus a number of samples, N_s, drawn in this fashion will asymptotically (as $N_s \rightarrow \infty$) converge to the true distribution [8], [9].

Using this sampling procedure we sample 50 sensors from the optimal distribution, as well as a uniform distribution, and calculate P_{SS} numerically for each placement. We repeat this 10,000 times in a Monte Carlo fashion. The result of this procedure is shown in Figure 3, where histograms of P_{SS} corresponding to each sampling of 50 sensors are shown, with results from the optimal distribution on top, and the uniform results on the bottom. It is seen that sample placements from the optimal distribution consistently outperform sample placements from the uniform distribution, as predicted. The expected coverage from each is found by taking the mean of these samples (P_{SS}), and is 0.306 for the optimal versus 0.2356 for the uniform.

We can use these results to gain guidance on sensor placement for this particular problem. For instance, if we take the random sample placements from figure 3 (optimal), corresponding to the maximum and minimum P_{SS}, we get the placements shown in figures 4 and 5, respectively. Focusing on the placement in figure 4, we see that the placement follows the distribution reasonably well, but where outliers are evident, they seem to occur in clusters of three. Since the multiple detection requirement is k = 3 for this case, this placement performs very well. Turning our attention to figure 5, which is the placement corresponding to the worst coverage, we see that the placement of the sensors is uneven, in that one section has sensors packed too densely with respect to the detection criteria, while another section has sensors that are too sparse. This leads to very poor overall field detection performance.

This example illustrates the utility of the optimization framework in providing insight to sensor placements which maximize track coverage for given target dynamics (stochastic) and sensor performance.

B. Example 2: Spatially Dependent Sensor Performance

In this example we use the same fusion parameters and target dynamics as before, but now we define a positional dependence on sensor performance within the search region. Figure 6 shows the positional dependence of R_D over the search region $\mathcal{S}$, which we use in this example. The detection radius takes values in the set, [50 100 200 300 400], as a function of radial distance from the origin (lower left corner). Including this in our sensor characterization and running the optimization procedure, we get the result shown in figure 7. The procedure was again started with equal weights (P_{SS} = 0.2499) and resulted in a P_{SS} of 0.3954. The numerical sampling procedure resulted in expected coverage (restricted

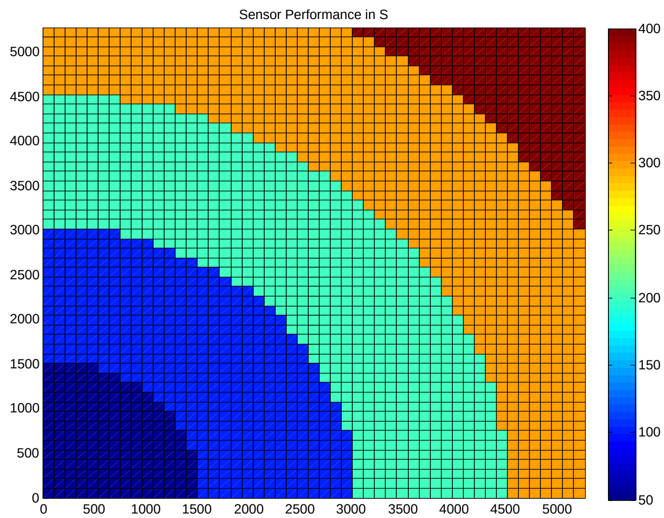

Fig. 6. Sensor Performance R_D as a function of position

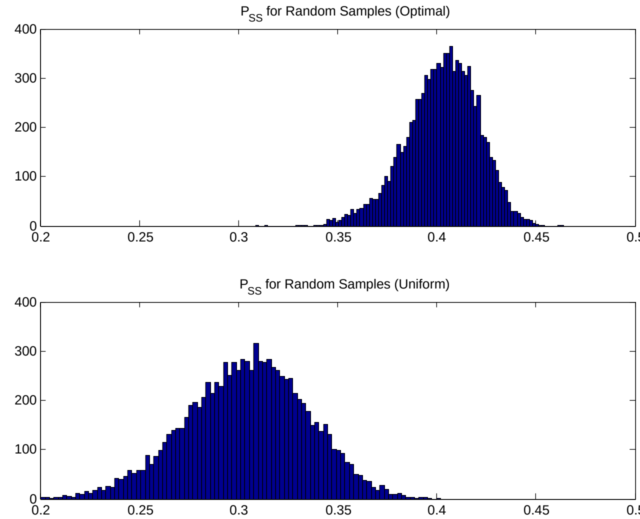

Fig. 8. Example 2 P_{SS} Sampling Results

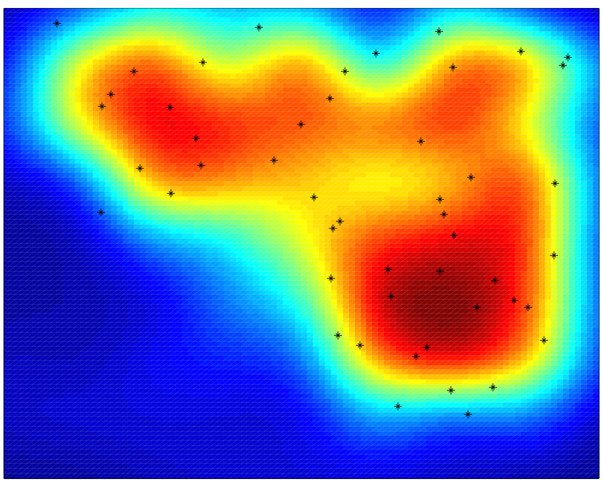

Fig. 7. Example 2 Genetic Result with Placement (max P_{SS})

to $\mathcal{S}$) of 0.4026 (optimal) versus 0.3035 (uniform) where histograms of P_{SS} are shown in figure 8. Once again we see that the optimization procedure results in better coverage than random placement.

The sensor distribution that is optimal for the search objective (figure 7) has the dominant mass where the sensors perform better, with very little mass placed where the sensors perform poorly. Again this is a consequence of the sparsity of the sensors available, combined with the fusion criteria. Included in figure 7 is the sensor placement corresponding to the random sample that performed best (out of the 10,000 realizations) with respect to P_{SS}. This placement is consistent with the underlying sensor distribution and resulted in a P_{SS} of 0.4638.

IV. Conclusion

In this paper we defined a search objective which emphasizes coverage of target tracks based on prior knowledge in an undersea surveillance problem. This objective was parameterized for use in a genetic optimization algorithm, to find a sensor distribution, which maximizes track coverage, for a fixed number of sensors. This optimization framework allowed us to explore the tradeoff of deploying sensors with respect to the optimal sensor distribution, versus randomly deploying sensors. Numerical examples showed that performance gains can be achieved, particularly when the number of sensors available only allow sparse coverage.

Acknowledgment

This work was supported by the Office of Naval Research Code 321MS and by the In-House Laboratory Independent Research Program of the Naval Undersea Warfare Center.

References

[1] R. Viswanathan and P. K. Varshney, *"Distributed Detection with Multiple Sensors: Part I – Fundamentals"*, Proceedings of the IEEE, Vol. 85, No. 1, January 1997.

[2] S. S. Dhillon and K. Chakrabarty, *"Sensor Placement for Effective Coverage and Surveillance in Distributed Sensor Networks"*, WCNC 2003, Volume 3, 16-20 March 2003 Page(s):1609 - 1614 vol.3.

[3] K. Chakrabarty, S. S. Iyengar, Q. Hairong, and C. Eungchun, *"Grid coverage for surveillance and target location in distributed sensor networks"* IEEE Transactions on Computers, Volume 51, Issue 12, Dec. 2002 Page(s):1448 - 1453.

[4] X. Zhang and S. B. Wicker, *"On the Optimal Distribution of Sensors in a Random Field"*, ACM Transactions on Sensor Networks, Vol. 1, No. 2, November 2005, Pages 301-306.

[5] T. A. Wettergren, *"Performance of search via track-before detect for distributed sensor networks"*, To Appear in IEEE Transactions on Aerospace and Electronic Systems.

[6] G. McLachlan and D. Peel, *"Finite Mixture Models"*, Wiley-Interscience, 2000.

[7] D. E.Goldberg, *"Genetic Algorithms"*, Addison Wesley Longman, 1999.

[8] K. L. Chung, *"A Course in Probability Theory"*, Academic Press, 2001.

[9] A. B. Owen, *"Empirical Likelihood"*, Monograph on Statistics and Applied Probability 92, Chapman Hall / CRC, 2001.